

WEST Search History

[Hide Items](#) | [Restore](#) | [Clear](#) | [Cancel](#)

DATE: Monday, April 17, 2006

<u>Hide?</u>	<u>Set Name</u>	<u>Query</u>	<u>Hit Count</u>
<i>DB=PGPB,USPT,USOC,EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=ADJ</i>			
<input type="checkbox"/>	L38	L37 and (disk near5 volume)	3
<input type="checkbox"/>	L37	(network\$3 and file\$1 and epoch and id\$) and @py<=2000	128
<input type="checkbox"/>	L36	L34 and disk\$1	5
<input type="checkbox"/>	L35	L34 and raid	0
<input type="checkbox"/>	L34	L32 and (logical near5 volume)	9
<input type="checkbox"/>	L33	L32 and epoch	0
<input type="checkbox"/>	L32	(6119131 or 6421684 or 6496839 or 6654881 or 6665786).pn.	10
<input type="checkbox"/>	L31	(data near5 structure) and (cluster\$1 near5 id\$) and (logical near5 volume) and (volume near5 id\$) and data and storage and raid and @py<=2000	0
<input type="checkbox"/>	L30	L29 and epoch	1
<input type="checkbox"/>	L29	L28 and server\$1 and configuration	18
<input type="checkbox"/>	L28	(logical near5 volume) and (compar\$3 or match\$3) and extent and (copy\$3 or transfer\$3) and memory and raid and cache and status and @py<=2000	22
<input type="checkbox"/>	L27	(raid near5 volume) and (raid near5 epoch)	0
<input type="checkbox"/>	L26	(epoch near5 id\$) and @py<=2000	4
<input type="checkbox"/>	L25	(epoch near5 id\$) and (raid near5 volume) and cluster\$3 and @py<=2000	0
<input type="checkbox"/>	L24	L23 and (compar\$3 near5 volume)	0
<input type="checkbox"/>	L23	L22 and (multiple near5 disk\$1)	5
<input type="checkbox"/>	L22	(raid and extent and size and logical and volume and configuration and disk\$1 and cluster\$3) and @py<2000	13
<input type="checkbox"/>	L21	L20 and timestamp\$3	5
<input type="checkbox"/>	L20	(disk\$1 same (logical near5 volume)) and configuration and status and compar\$3 and (index\$3 or custer\$3) and extent and size and @py<=2000	22
<input type="checkbox"/>	L19	L18 and logical	3
<input type="checkbox"/>	L18	L17 and cluster	3
<input type="checkbox"/>	L17	(raid and epoch) and @py<=2000	13
<input type="checkbox"/>	L16	(disk near5 epoch) and @py<=2000	2
<input type="checkbox"/>	L15	L14 and epoch	0
<input type="checkbox"/>	L14	L13 and (logical near5 volume)	14
<input type="checkbox"/>	L13	L10 and cluster\$1	63
<input type="checkbox"/>	L12	L10 and (epoch and configuration and timestamp\$3)	3

<input type="checkbox"/>	L11	L10 and (epoch near5 identifier\$1)	0
<input type="checkbox"/>	L10	(disk\$1 and extent and volume and mirror\$3 and block\$1 and logical and remote and file\$1) and @py<=2000	188
<input type="checkbox"/>	L9	(logical and file\$1 and disk and raid and logical and identifier\$1 and extent and epoch) and @py<=2000	5
<input type="checkbox"/>	L8	(epoch and disk and volume and raid and logical) and @py<=2000	6
<input type="checkbox"/>	L7	L6 and epoch	1
<input type="checkbox"/>	L6	L5 and (disk near5 block\$1)	6
<input type="checkbox"/>	L5	L4 and (disk near5 extent\$1)	9
<input type="checkbox"/>	L4	(raid and disk\$1 and logical and volume and extent\$1 and configuration and mapp\$3) and @py<=2000	52
<input type="checkbox"/>	L3	L2 and epoch	3
<input type="checkbox"/>	L2	(raid and mirror\$3 and volume and logical and block\$1 and mapp\$3 and extent\$1) and @py<=2000	37
<input type="checkbox"/>	L1	(disk and raid and identifier\$1 and logical and volume and configuration and data and structgure and extent) and @py<=1999	0

END OF SEARCH HISTORY

WEST Search History

DATE: Monday, April 17, 2006

<u>Hide?</u>	<u>Set</u>	<u>Name</u>	<u>Query</u>	<u>Hit Count</u>
			<i>DB=PGPB,USPT,USOC,EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=ADJ</i>	
<input type="checkbox"/>	L13	L12	and (copy near5 epoch)	0
<input type="checkbox"/>	L12	(read near5 epoch)	and @py<=2000	25
<input type="checkbox"/>	L11	L10	and read\$3	12
<input type="checkbox"/>	L10	L9	and epoch	13
<input type="checkbox"/>	L9	L8	and (logical near5 volume)	673
<input type="checkbox"/>	L8	L7	and cluster\$3	20839
<input type="checkbox"/>	L7	(disk and volume and cluster\$3)		20839
<input type="checkbox"/>	L6	L5	and (mirror\$3 near5 volume)	5
<input type="checkbox"/>	L5	disk	and epoch and read\$3 and copy and value\$1 and data and structure and volume and cluster\$3	96
<input type="checkbox"/>	L4	(disk near5 epoch)	and (read\$3 near5 epoch)	4
<input type="checkbox"/>	L3	(read\$3 and epoch)	and value\$1 and mirroed and copy and logical and volume and disk)	0
<input type="checkbox"/>	L2	(read\$3 and epoch)	and value\$1 and mirroed and copy and logical and volume and raid)	0
<input type="checkbox"/>	L1	(read\$3 and epoch)	and value\$1 and mirroed and copy and logical and volume and raid) and @py<=2000	0

END OF SEARCH HISTORY



[Home](#) | [Login](#) | [Logout](#) | [Access Information](#) | [Alerts](#) |
Welcome United States Patent and Trademark Office

[Search Results](#)

[BROWSE](#)

[SEARCH](#)

[IEEE XPLOR GUIDE](#)

Results for "(logical<in>metadata) <and> (cluster<in>metadata) <and> (disk<in>metadata)"
 Your search matched 3 of 1340257 documents.

[e-mail](#)

A maximum of 100 results are displayed, 25 to a page, sorted by **Relevance in Descending order**.

» [Search Options](#)

[View Session History](#)

[Modify Search](#)

[New Search](#)

((logical<in>metadata) <and> (cluster<in>metadata) <and> (disk<in>metadata))

Check to search only within this results set

Display Format: Citation Citation & Abstract

[Select All](#) [Deselect All](#)

» [Key](#)

IEEE JNL IEEE Journal or Magazine

IEE JNL IEE Journal or Magazine

IEEE CNF IEEE Conference Proceeding

IEE CNF IEE Conference Proceeding

IEEE STD IEEE Standard

1. **Back to the future: dynamic hierarchical clustering**

Chendong Zou; Salzberg, B.; Ladin, R.;

[Data Engineering, 1998. Proceedings., 14th International Conference on](#)
 23-27 Feb. 1998 Page(s):578 - 587

Digital Object Identifier 10.1109/ICDE.1998.655821

[AbstractPlus](#) | Full Text: [PDF\(1420 KB\)](#) [IEEE CNF](#)
[Rights and Permissions](#)

2. **Logical vs. physical disk shadowing**

Graefe, G.; Shapiro, L.D.;

[Applied Computing, 1991.. \[Proceedings of the 1991\] Symposium on](#)
 3-5 April 1991 Page(s):91

Digital Object Identifier 10.1109/SOAC.1991.143853

[AbstractPlus](#) | Full Text: [PDF\(96 KB\)](#) [IEEE CNF](#)
[Rights and Permissions](#)

3. **The virtual cluster: a dynamic network environment for exploitation of idl**

De Rose, C.; Blanco, F.; Maillard, N.; Saikoski, K.; Novaes, R.; Richard, O.; Ri
[Computer Architecture and High Performance Computing, 2002. Proceedings.](#)
 on

28-30 Oct. 2002 Page(s):141 - 148

Digital Object Identifier 10.1109/CAHPC.2002.1180770

[AbstractPlus](#) | Full Text: [PDF\(426 KB\)](#) [IEEE CNF](#)
[Rights and Permissions](#)

[Help](#) [Contact Us](#) [Privacy &](#)

© Copyright 2006 IEEE -

Indexed by
 Inspec

[Home](#) | [Login](#) | [Logout](#) | [Access Information](#) | [Alerts](#) |

Welcome United States Patent and Trademark Office

 [Search Results](#)[BROWSE](#)[SEARCH](#)[IEEE Xplore GUIDE](#) [e-mail](#)

Results for "(logical<in>metadata) <and> (volume<in>metadata) <and> (extent<in>meta...)"

Your search matched 2 of 1340257 documents.

A maximum of 100 results are displayed; 25 to a page, sorted by **Relevance in Descending order**.» [Search Options](#)[View Session History](#)[Modify Search](#)[New Search](#)» [Key](#)

IEEE JNL IEEE Journal or Magazine

IEE JNL IEE Journal or Magazine

IEEE CNF IEEE Conference Proceeding

IEE CNF IEE Conference Proceeding

IEEE STD IEEE Standard

 [Select All](#) [Deselect All](#) Check to search only within this results setDisplay Format: Citation Citation & Abstract

1. Postflood damage evaluation using Landsat TM and ETM+ data integrated
Gianinetto, M.; Villa, P.; Lechi, G.;
[Geoscience and Remote Sensing, IEEE Transactions on](#)
Volume 44, Issue 1, Jan. 2006 Page(s):236 - 243
Digital Object Identifier 10.1109/TGRS.2005.859952

[AbstractPlus](#) | Full Text: [PDF\(3368 KB\)](#) [IEEE JNL](#)
[Rights and Permissions](#)

2. Digital Microelectronic Equipment Problems and Potentials

Coleman, A.;
[Product Engineering and Production, IRE Transactions on](#)
Volume 5, Issue 2, Jun 1961 Page(s):3 - 12

[AbstractPlus](#) | Full Text: [PDF\(1168 KB\)](#) [IEEE JNL](#)
[Rights and Permissions](#)[Help](#) [Contact Us](#) [Privacy &](#)

© Copyright 2006 IEEE -

Indexed by
Inspec


[Home](#) | [Login](#) | [Logout](#) | [Access Information](#) | [Alerts](#) |

Welcome United States Patent and Trademark Office

 [Search Results](#)[BROWSE](#)[SEARCH](#)[IEEE Xplore GUIDE](#)

Results for "(logical<in>metadata) <and> (volume<in>metadata) <and> (disk<in>metadata)"

 [e-mail](#)

Your search matched 7 of 1340257 documents.

A maximum of 100 results are displayed, 25 to a page, sorted by **Relevance in Descending order**.» **Search Options**[View Session History](#)**Modify Search**[New Search](#) Check to search only within this results set» **Key**Display Format: Citation Citation & Abstract**IEEE JNL** IEEE Journal or Magazine [Select All](#) [Deselect All](#)**IEE JNL** IEE Journal or Magazine**IEEE CNF** IEEE Conference Proceeding**IEE CNF** IEE Conference Proceeding**IEEE STD** IEEE Standard

1. **A Block-level Security based on Hierarchical Logical Volume of Fibre Channel**
 Ke Zhou; Feng, D.; Wang, F.; Zhan Shi;
Embedded Software and Systems, 2005. Second International Conference on
16-18 Dec. 2005 Page(s):574 - 577
 Digital Object Identifier 10.1109/ICESS.2005.1

[AbstractPlus](#) | [Full Text: PDF\(176 KB\)](#) [IEEE CNF](#)
[Rights and Permissions](#)

2. **Performance of a storage system for supporting different video types and standards**
 Jonathan Chien-Liang Liu; Jenwei Hsieh; Du, D.H.C.; Meng-Jou Lin;
Selected Areas in Communications, IEEE Journal on
Volume 14, Issue 7, Sept. 1996 Page(s):1314 - 1331
 Digital Object Identifier 10.1109/49.536482

[AbstractPlus](#) | [References](#) | [Full Text: PDF\(1904 KB\)](#) [IEEE JNL](#)
[Rights and Permissions](#)

3. **Performance of a storage system for supporting different video types and standards**
 Liu, J.C.L.; Hsieh, J.; Du, D.H.C.; Mengjou Lin;
INFOCOM '96. Fifteenth Annual Joint Conference of the IEEE Computer Society on the Next Generation. Proceedings IEEE
Volume 1, 24-28 March 1996 Page(s):2 - 9 vol.1
 Digital Object Identifier 10.1109/INFCOM.1996.497871

[AbstractPlus](#) | [Full Text: PDF\(780 KB\)](#) [IEEE CNF](#)
[Rights and Permissions](#)

4. **Content protection technology for a novel removable drive**
 Hirai, T.;

Magnetics, IEEE Transactions on
Volume 41, Issue 2, Feb. 2005 Page(s):860 - 869
 Digital Object Identifier 10.1109/TMAG.2004.840298[AbstractPlus](#) | [Full Text: PDF\(832 KB\)](#) [IEEE JNL](#)
[Rights and Permissions](#)

5. **CD-ROM system based on the NEWS workstation**
 Kimura, K.; Demura, A.; Igarashi, T.; Takeyari, Y.;

Compcon Spring '88. Thirty-Third IEEE Computer Society International Conference on
Papers
 29 Feb.-3 March 1988 Page(s):274 - 276

Digital Object Identifier 10.1109/CMPCON.1988.4873

[AbstractPlus](#) | Full Text: [PDF\(200 KB\)](#) IEEE CNF
[Rights and Permissions](#)

- 6. **An overview of optical disk standards**
Williams, D.I.;
[Storage and Recording Systems, 1994., International Conference on](#)
5-7 Apr 1994 Page(s):134 - 136
[AbstractPlus](#) | Full Text: [PDF\(212 KB\)](#) IEE CNF

- 7. **CLARE-a CLAuse Retrieval Engine for large Prolog databases**
Kam-Fai Wong; Williams, M.H.;
[VLSI and Architectures for Symbolic Processing, IEE Colloquium on](#)
9 Mar 1989 Page(s):3/1 - 3/3
[AbstractPlus](#) | Full Text: [PDF\(164 KB\)](#) IEE CNF

[Help](#) [Contact Us](#) [Privacy &](#)

© Copyright 2006 IEEE -

Indexed by
 Inspec®


[Home](#) | [Login](#) | [Logout](#) | [Access Information](#) | [Alerts](#) |

Welcome United States Patent and Trademark Office

 Search Results**BROWSE****SEARCH****IEEE Xplore GUIDE**

Results for "(logical<in>metadata) <and> (volume<in>metadata) <and> (comparing<..."

 e-mailYour search matched **6** of **1340257** documents.A maximum of **100** results are displayed, **25** to a page, sorted by **Relevance** in **Descending** order.» **Search Options**[View Session History](#)[Modify Search](#)[New Search](#) Check to search only within this results set» **Key**Display Format: Citation Citation & Abstract**IEEE JNL** IEEE Journal or Magazine [Select All](#) [Deselect All](#)**IEE JNL** IEE Journal or Magazine**IEEE CNF** IEEE Conference Proceeding**IEE CNF** IEE Conference Proceeding**IEEE STD** IEEE Standard

1. **Influence of Adaptive Data Layouts on Performance in Dynamically Changing Environments**

Brinkmann, A.; Effert, S.; Heidebuer, M.; Vodisek, M.; [Parallel, Distributed, and Network-Based Processing, 2006. PDP 2006. 14th E](#)
[International Conference on](#)

15-17 Feb. 2006 Page(s):155 - 162

Digital Object Identifier 10.1109/PDP.2006.44

[AbstractPlus](#) | Full Text: [PDF\(280 KB\)](#) [IEEE CNF](#)[Rights and Permissions](#)

2. **Enabling technologies for integrated system-on-a-package for the next generation aerospace applications**

Young, J.J.; Malshe, A.P.; Brown, W.D.; Lenihan, T.; Albert, D.; Ozguz, V.; [Aerospace Conference Proceedings, 2002. IEEE](#)

Volume 5, 2002 Page(s):5-2177 - 5-2184 vol.5

Digital Object Identifier 10.1109/AERO.2002.1035384

[AbstractPlus](#) | Full Text: [PDF\(979 KB\)](#) [IEEE CNF](#)[Rights and Permissions](#)

3. **A survey of available underwater electric propulsion technologies and implications for platform system safety**

Winchester, C.; Govar, J.; Banner, J.; Squires, T.; Smith, P.; [Autonomous Underwater Vehicles, 2002. Proceedings of the 2002 Workshop on](#)

20-21 June 2002 Page(s):129 - 135

Digital Object Identifier 10.1109/AUV.2002.1177215

[AbstractPlus](#) | Full Text: [PDF\(592 KB\)](#) [IEEE CNF](#)[Rights and Permissions](#)

4. **A comparative study of genetic sequence classification algorithms**

Mukhopadhyay, S.; Changhong Tang; Huang, J.; Mulong Yu; Palakal, M.; [Neural Networks for Signal Processing, 2002. Proceedings of the 2002 12th IEEE](#)

4-6 Sept. 2002 Page(s):57 - 66

Digital Object Identifier 10.1109/NNSP.2002.1030017

[AbstractPlus](#) | Full Text: [PDF\(459 KB\)](#) [IEEE CNF](#)[Rights and Permissions](#)

5. **Modeling human left ventricle and aortic function using VisSim**

Johnson, B.A.; Summers, R.L.; Cathcart, T.P.;

[Biomedical Engineering Conference, 1997., Proceedings of the 1997 Sixteenth](#)
4-6 April 1997 Page(s):374 - 377
Digital Object Identifier 10.1109/SBEC.1997.583315
[AbstractPlus](#) | Full Text: [PDF\(324 KB\)](#) IEEE CNF
[Rights and Permissions](#)



6. The design of efficient XML document model

Jun Wen; Rui Zhang; Xianliang Lu;
[Machine Learning and Cybernetics, 2002. Proceedings, 2002 International Co](#)
Volume 2, 4-5 Nov. 2002 Page(s):1102 - 1106 vol.2
Digital Object Identifier 10.1109/ICMLC.2002.1174555
[AbstractPlus](#) | Full Text: [PDF\(571 KB\)](#) IEEE CNF
[Rights and Permissions](#)

[Help](#) [Contact Us](#) [Privacy &](#)

© Copyright 2006 IEEE -

Indexed by
Inspec

[Home](#) | [Login](#) | [Logout](#) | [Access Information](#) | [Alerts](#) |

Welcome United States Patent and Trademark Office

 [Search Results](#)[BROWSE](#)[SEARCH](#)[IEEE XPLOR GUIDE](#)

Results for "(logical<in>metadata) <and> (volume<in>metadata) <and> (epoch<in>meta...)"

 [e-mail](#)

Your search matched 1 of 1340257 documents.

A maximum of 100 results are displayed, 25 to a page, sorted by **Relevance** in **Descending** order.» [Search Options](#)[View Session History](#)[Modify Search](#)[New Search](#)» [Key](#)

IEEE JNL IEEE Journal or Magazine

IEE JNL IEE Journal or Magazine

IEEE CNF IEEE Conference Proceeding

IEE CNF IEE Conference Proceeding

IEEE STD IEEE Standard

Display Format: Citation Citation & Abstract

1. **Studies on parallel and distributed RS image issuance system based on International**
Wu, H.Q.; Chi, T.H.; Fang, J.Y.; Zhang, X.;
Geoscience and Remote Sensing Symposium, 2003. IGARSS '03. Proceeding
Volume 6, 21-25 July 2003 Page(s):3790 - 3792 vol.6
Digital Object Identifier 10.1109/IGARSS.2003.1295271
[AbstractPlus](#) | Full Text: [PDF\(1314 KB\)](#) [IEEE CNF Rights and Permissions](#)

[Help](#) [Contact Us](#) [Privacy &](#)

© Copyright 2006 IEEE –

Indexed by



Home | Login | Logout | Access Information | Alerts |
Welcome United States Patent and Trademark Office

Search Session History

BROWSE

SEARCH

IEEE XPLORER GUIDE

Mon, 17 Apr 2006, 9:35:35 AM EST

Edit an existing query or
compose a new query in the
Search Query Display.

Search Query Display

Select a search number (#)
to:

- Add a query to the Search Query Display
- Combine search queries using AND, OR, or NOT
- Delete a search
- Run a search

Recent Search Queries

- #1 ((disk<in>metadata) <and> (epoch<in>metadata)<and>
(identifiers<in>metadata)
- #2 ((disk<in>metadata) <and> (epoch<in>metadata)<and>
(id<in>metadata)
- #3 ((disk<in>metadata) <and> (epoch<in>metadata)<and>
(logical<in>metadata)
- #4 ((disk<in>metadata) <and> (epoch<in>metadata)<and>
(cluster<in>metadata)
- #5 ((raid<in>metadata) <and> (epoch<in>metadata)<and>
(cluster<in>metadata)
- #6 ((raid<in>metadata) <and> (volume<in>metadata)<and>
(cluster<in>metadata)
- #7 ((raid<in>metadata) <and> (volume<in>metadata)<and>
(cluster<in>metadata)
- #8 ((logical<in>metadata) <and> (cluster<in>metadata)<and>
(disk<in>metadata)
- #9 ((logical<in>metadata) <and> (cluster<in>metadata)<and>
(disk<in>metadata)
- #10 ((logical<in>metadata) <and> (volume<in>metadata)<and>
(mirror<in>metadata)
- #11 ((logical<in>metadata) <and> (volume<in>metadata)<and>
(backup<in>metadata)
- #12 ((logical<in>metadata) <and> (volume<in>metadata)<and>
(extent<in>metadata)
- #13 ((logical<in>metadata) <and> (volume<in>metadata)<and>
(extent<in>metadata)
- #14 ((logical<in>metadata) <and> (volume<in>metadata)<and>
(disk<in>metadata)
- #15 ((logical<in>metadata) <and> (volume<in>metadata)<and>
(comparing<in>metadata)

#16 ((logical<in>metadata) <and> (volume<in>metadata))<and>
(epoch<in>metadata)

[Search Results]

[Help](#) [Contact Us](#) [Privacy &](#)

© Copyright 2006 IEEE -

Indexed by
Inspec



[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

Search: The ACM Digital Library The Guide

raid epoch identifiers

SEARCH

THE ACM DIGITAL LIBRARY

[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used [raid epoch identifiers](#)

Found 336 of 175,083

Sort results by

relevance

[Save results to a Binder](#)

Try an [Advanced Search](#)

Display results

expanded form

[Search Tips](#)

Try this search in [The ACM Guide](#)

Open results in a new window

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale

1 Petal: distributed virtual disks

Edward K. Lee, Chandramohan A. Thekkath

September 1996 **ACM SIGPLAN Notices , ACM SIGOPS Operating Systems Review , Proceedings of the seventh international conference on Architectural support for programming languages and operating systems ASPLOS-VII**, Volume 31 , 30 Issue 9 , 5

Publisher: ACM Press

Full text available: [pdf\(1.10 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The ideal storage system is globally accessible, always available, provides unlimited performance and capacity for a large number of clients, and requires no management. This paper describes the design, implementation, and performance of Petal, a system that attempts to approximate this ideal in practice through a novel combination of features. Petal consists of a collection of network-connected servers that cooperatively manage a pool of physical disks. To a Petal client, this collection appear ...

2 SPV: secure path vector routing for securing BGP

Yih-Chun Hu, Adrian Perrig, Marvin Sirbu

August 2004 **ACM SIGCOMM Computer Communication Review , Proceedings of the 2004 conference on Applications, technologies, architectures, and protocols for computer communications SIGCOMM '04**, Volume 34 Issue 4

Publisher: ACM Press

Full text available: [pdf\(236.82 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

As our economy and critical infrastructure increasingly relies on the Internet, the insecurity of the underlying border gateway routing protocol (BGP) stands out as the Achilles heel. Recent misconfigurations and attacks have demonstrated the brittleness of BGP. Securing BGP has become a priority. In this paper, we focus on a viable deployment path to secure BGP. We analyze security requirements, and consider tradeoffs of mechanisms that achieve the requirements. In particular, we study how to se ...

Keywords: BGP, Border Gateway Protocol, interdomain routing, routing, security

3 Improving storage system availability with D-GRAID

Muthian Sivathanu, Vijayan Prabhakaran, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau

May 2005 **ACM Transactions on Storage (TOS)**, Volume 1 Issue 2

Publisher: ACM Press

Full text available: [pdf\(700.30 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

We present the design, implementation, and evaluation of D-GRAID, a gracefully degrading and quickly recovering RAID storage array. D-GRAID ensures that most files within the file system remain available even when an unexpectedly high number of faults occur. D-GRAID achieves high availability through aggressive replication of semantically critical data, and fault-isolated placement of logically related data. D-GRAID also recovers from failures quickly, restoring only live file system data to a h ...

Keywords: Block-based storage, Disk array, RAID, fault isolation, file systems, smart disks

4 Manageability, availability, and performance in porcupine: a highly scalable, cluster-based mail service

 Yasushi Saito, Brian N. Bershad, Henry M. Levy

August 2000 **ACM Transactions on Computer Systems (TOCS)**, Volume 18 Issue 3

Publisher: ACM Press

Full text available: [pdf\(2.52 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

This paper describes the motivation, design and performance of Porcupine, a scalable mail server. The goal of Porcupine is to provide a highly available and scalable electronic mail service using a large cluster of commodity PCs. We designed Porcupine to be easy to manage by emphasizing dynamic load balancing, automatic configuration, and graceful degradation in the presence of failures. Key to the system's manageability, availability, and performance is that sessions, data, and underlying ...

Keywords: cluster, distributed systems, email, group membership protocol, load balancing, replication

5 Lightweight recoverable virtual memory

 M. Satyanarayanan, Henry H. Mashburn, Puneet Kumar, David C. Steere, James J. Kistler
December 1993 **ACM SIGOPS Operating Systems Review , Proceedings of the fourteenth ACM symposium on Operating systems principles SOSP '93**, Volume 27 Issue 5

Publisher: ACM Press

Full text available: [pdf\(1.53 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Recoverable virtual memory refers to regions of a virtual address space on which transactional guarantees are offered. This paper describes RVM, an efficient, portable, and easily used implementation of recoverable virtual memory for Unix environments. A unique characteristic of RVM is that it allows independent control over the transactional properties of atomicity, permanence, and serializability. This leads to considerable flexibility in the use of RVM, potentially enlarging the ...

6 Lightweight recoverable virtual memory

 M. Satyanarayanan, Henry H. Mashburn, Puneet Kumar, David C. Steere, James J. Kistler
February 1994 **ACM Transactions on Computer Systems (TOCS)**, Volume 12 Issue 1

Publisher: ACM Press

Full text available: [pdf\(1.73 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Recoverable virtual memory refers to regions of a virtual address space on which

transactional guarantees are offered. This article describes RVM, an efficient, portable, and easily used implementation of recoverable virtual memory for Unix environments. A unique characteristic of RVM is that it allows independent control over the transactional properties of atomicity, permanence, and serializability. This leads to considerable flexibility in the use of RVM, potentially enabling ...

Keywords: Camelot, Coda, RVM, Unix, logging, paging, persistence, scalability, throughput, truncation

7 The monitoring and early detection of internet worms

Cliff C. Zou, Weibo Gong, Don Towsley, Lixin Gao

October 2005 **IEEE/ACM Transactions on Networking (TON)**, Volume 13 Issue 5

Publisher: IEEE Press

Full text available:  [pdf\(594.79 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

After many Internet-scale worm incidents in recent years, it is clear that a simple self-propagating worm can quickly spread across the Internet and cause severe damage to our society. Facing this great security threat, we need to build an early detection system that can detect the presence of a worm in the Internet as quickly as possible in order to give people accurate early warning information and possible reaction time for counteractions. This paper first presents an Internet worm monitoring ...

Keywords: computer network security, early detection, internet worm, network monitoring

8 Hibernator: helping disk arrays sleep through the winter



Qingbo Zhu, Zhifeng Chen, Lin Tan, Yuanyuan Zhou, Kimberly Keeton, John Wilkes

October 2005 **ACM SIGOPS Operating Systems Review , Proceedings of the twentieth ACM symposium on Operating systems principles SOSP '05**, Volume 39 Issue 5

Publisher: ACM Press

Full text available:  [pdf\(654.56 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Energy consumption has become an important issue in high-end data centers, and disk arrays are one of the largest energy consumers within them. Although several attempts have been made to improve disk array energy management, the existing solutions either provide little energy savings or significantly degrade performance for data center workloads. Our solution, Hibernator, is a disk array energy management system that provides improved energy savings while meeting performance goals. Hibernator co ...

Keywords: disk array, disk layout, energy management, performance guarantee, storage system

9 Distributed logging for transaction processing



Dean S. Daniels, Alfred Z. Spector, Dean S. Thompson

December 1987 **ACM SIGMOD Record , Proceedings of the 1987 ACM SIGMOD international conference on Management of data SIGMOD '87**, Volume 16 Issue 3

Publisher: ACM Press

Full text available:  [pdf\(1.51 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Increased interest in using workstations and small processors for distributed transaction processing raises the question of how to implement the logs needed for transaction recovery. Although logs can be implemented with data written to duplexed disks on each

processing node, this paper argues there are advantages if log data is written to multiple log server nodes. A simple analysis of expected logging loads leads to the conclusion that a high performance, microprocessor b ...

10 Gossip-based aggregation in large dynamic networks

 Márk Jelasity, Alberto Montresor, Ozalp Babaoglu
August 2005 **ACM Transactions on Computer Systems (TOCS)**, Volume 23 Issue 3

Publisher: ACM Press

Full text available:  pdf(533.89 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

As computer networks increase in size, become more heterogeneous and span greater geographic distances, applications must be designed to cope with the very large scale, poor reliability, and often, with the extreme dynamism of the underlying network.

Aggregation is a key functional building block for such applications: it refers to a set of functions that provide components of a distributed system access to global information including network size, average load, average uptime, location ...

Keywords: Gossip-based protocols, proactive aggregation

11 The Zebra striped network file system

 John H. Hartman, John K. Ousterhout
August 1995 **ACM Transactions on Computer Systems (TOCS)**, Volume 13 Issue 3

Publisher: ACM Press

Full text available:  pdf(2.76 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Zebra is a network file system that increases throughput by striping the file data across multiple servers. Rather than striping each file separately, Zebra forms all the new data from each client into a single stream, which it then stripes using an approach similar to a log-structured file system. This provides high performance for writes of small files as well as for reads and writes of large files. Zebra also writes parity information in each stripe in the style of RAID disk arrays; this ...

Keywords: RAID, log-based striping, log-structured file system, parity computation

12 The TickerTAIP parallel RAID architecture

 Pei Cao, Swee Boon Lin, Shivakumar Venkataraman, John Wilkes
August 1994 **ACM Transactions on Computer Systems (TOCS)**, Volume 12 Issue 3

Publisher: ACM Press

Full text available:  pdf(2.04 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Traditional disk arrays have a centralized architecture, with a single controller through which all requests flow. Such a controller is a single point of failure, and its performance limits the maximum number of disks to which the array can scale. We describe TickerTAIP, a parallel architecture for disk arrays that distributes the controller functions across several loosely coupled processors. The result is better scalability, fault tolerance, and flexibility. This article present ...

Keywords: RAID disk array, decentralized parity calculation, disk scheduling, distributed controller, fault tolerance, parallel controller, performance simulation

13

Epochs, configuration schema, and version cursors in the KBSA framework CCM

 **model**

John Kimball, Aaron Larson

May 1991 **Proceedings of the 3rd international workshop on Software configuration management**

Publisher: ACM Press

Full text available:  [pdf\(1.06 MB\)](#) Additional Information: [full citation](#), [references](#), [index terms](#)

14 The TickerTAIP parallel RAID architecture  Pei Cao, Swee Boon Lim, Shivakumar Venkataraman, John Wilkes

May 1993 **ACM SIGARCH Computer Architecture News , Proceedings of the 20th annual international symposium on Computer architecture ISCA '93**, Volume 21 Issue 2

Publisher: ACM Press

Full text available:  [pdf\(1.19 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Traditional disk arrays have a centralized architecture, with a single controller through which all requests flow. Such a controller is a single point of failure, and its performance limits the maximum size that the array can grow to. We describe here TickerTAIP, a parallel architecture for disk arrays that distributed the controller functions across several loosely-coupled processors. The result is better scalability, fault tolerance, and flexibility. This paper presents the Tic ...

15 Implementing sequentially consistent shared objects using broadcast and point-to-point communication  Alan Fekete, M. Frans Kaashoek, Nancy Lynch

January 1998 **Journal of the ACM (JACM)**, Volume 45 Issue 1

Publisher: ACM Press

Full text available:  [pdf\(257.13 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper presents and proves correct a distributed algorithm that implements a sequentially consistent collection of shared read/update objects. This algorithm is a generalization of one used in the Orca shared object system. The algorithm caches objects in the local memory of processors according to application needs; each read operation accesses a single copy of the object, while each update accesses all copies. The algorithm uses broadcast communication when it sends messages to replic ...

Keywords: Orca programming language, distributed shared memory, formal methods, input/output automata, ordered multicast, replicated data

16 Tolerating multiple failures in RAID architectures with optimal storage and uniform declustering  Guillermo A. Alvarez, Walter A. Burkhard, Flaviu Cristian

May 1997 **ACM SIGARCH Computer Architecture News , Proceedings of the 24th annual international symposium on Computer architecture ISCA '97**, Volume 25 Issue 2

Publisher: ACM Press

Full text available:  [pdf\(1.50 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

We present DATUM, a novel method for tolerating multiple disk failures in disk arrays. DATUM is the first known method that can mask any given number of failures, requires an optimal amount of redundant storage space, and spreads reconstruction accesses uniformly over disks in the presence of failures without needing large layout tables in

controller memory. Our approach is based on information dispersal, a coding technique that admits an efficient hardware implementation. As t ...

17 The Zebra striped network file system

 John H. Hartman, John K. Ousterhout

December 1993 **ACM SIGOPS Operating Systems Review , Proceedings of the fourteenth ACM symposium on Operating systems principles SOSP '93**, Volume 27 Issue 5

Publisher: ACM Press

Full text available:  pdf(1.93 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Zebra is a network file system that increases throughput by striping file data across multiple servers. Rather than striping each file separately, Zebra forms all the new data from each client into a single stream, which it then stripes using an approach similar to a log-structured file system. This provides high performance for writes of small files as well as for reads and writes of large files. Zebra also writes parity information in each stripe in the style of RAID disk arrays; this increase ...

18 A 50-Gb/s IP router

Craig Partridge, Philip P. Carvey, Ed Burgess, Isidro Castineyra, Tom Clarke, Lise Graham, Michael Hathaway, Phil Herman, Allen King, Steve Kohalmi, Tracy Ma, John Mcallen, Trevor Mendez, Walter C. Milliken, Ronald Pettyjohn, John Rokosz, Joshua Seeger, Michael Sollins, Steve Storch, Benjamin Tober, Gregory D. Troxel

June 1998 **IEEE/ACM Transactions on Networking (TON)**, Volume 6 Issue 3

Publisher: IEEE Press

Full text available:  pdf(133.28 KB)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#), [review](#)

Keywords: data communications, internetworking, packet switching, routing

19 Guide for the use of the Ada Ravenscar Profile in high integrity systems

 Alan Burns, Brian Dobbing, Tullio Vardanega

June 2004 **ACM SIGAda Ada Letters**, Volume XXIV Issue 2

Publisher: ACM Press

Full text available:  pdf(548.17 KB)

Additional Information: [full citation](#), [references](#)

20 Routing: Ensuring cache freshness in on-demand ad hoc network routing protocols

 Yih-Chun Hu, David B. Johnson

October 2002 **Proceedings of the second ACM international workshop on Principles of mobile computing**

Publisher: ACM Press

Full text available:  pdf(131.62 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In a wireless ad hoc network, nodes cooperate to forward packets for each other over possibly multi-hop paths, to allow nodes not within direct wireless transmission range to communicate. Many routing protocols have been proposed for the ad hoc network environment, several of which operate on-demand and utilize a *route cache* listing links that this node has learned. In such protocols, aggressive caching of overheard routes can significantly improve performance; in particular, overhead can ...

Keywords: DSR, Dynamic Source Routing, ad hoc networks, bounded latency, epoch numbers, route cache, theory

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)

 [Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

Search: The ACM Digital Library The Guide

extent ekepoch raid virtual disks

THE ACM DIGITAL LIBRARY

 [Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used [extent ekepoch raid virtual disks](#)

Found 13,572 of 175,083

Sort results by relevance Save results to a Binder
 Search Tips

Display results expanded form Open results in a new window

Try an [Advanced Search](#)
 Try this search in [The ACM Guide](#)

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale **1 Hierarchical disk cache management in RAID 5 controller**

Jung-ho Huh, Tae-mu Chang

December 2003 **Journal of Computing Sciences in Colleges**, Volume 19 Issue 2**Publisher:** Consortium for Computing Sciences in CollegesFull text available:  [pdf\(137.71 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

In RAID system, disk cache is one of the important elements in improving the system performance. Two-level cache displays superior performance in comparison to single cache and is effective in temporal and spatial locality. The proposed cache system consists in two levels. The first level cache is a set associative cache with small block size whereas the second level cache is a fully associative cache with large block size. In this paper, a RAID 5 disk cache model is presented that is located on ...

2 Data partitioning and load balancing in parallel disk systems

Peter Scheuermann, Gerhard Weikum, Peter Zabback

February 1998 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 7 Issue 1**Publisher:** Springer-Verlag New York, Inc.Full text available:  [pdf\(310.27 KB\)](#) Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

Parallel disk systems provide opportunities for exploiting I/O parallelism in two possible ways, namely via inter-request and intra-request parallelism. In this paper, we discuss the main issues in performance tuning of such systems, namely striping and load balancing, and show their relationship to response time and throughput. We outline the main components of an intelligent, self-reliant file system that aims to optimize striping by taking into account the requirements of the applications, an ...

Keywords: Data allocation, Disk cooling, File striping, Load balancing, Parallel disk systems, Performance tuning

3 Technology to enable learning: Strategic decisions on technology selections for facilitating a network/systems laboratory using real options & total cost of ownership theories

Kmfong Lei, Phillip T. Rawles

October 2003 **Proceedings of the 4th conference on Information technology curriculum CITEC4 '03****Publisher:** ACM Press

Full text available:  pdf(407.50 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

This paper addresses the selection of technologies that provide each student group a dedicated environment on a non-dedicated host machine. The authors investigated different combinations of enabling technologies and approaches, such as virtual machine technology, storage technology, and host operating system. Performance tests were developed and executed to profile the performance of the technologies. The results of this work provide an evaluation of the studied technologies and a selection gui ...

Keywords: VMware, course development, curriculum, end-user computing, innovative lab strategies in IT, interesting applications in IT, networking, operating systems, systems software

4 False sharing problems in cluster-based disk arrays

 Hai Jin, Kai Hwang

February 1999 **Proceedings of the 1999 ACM symposium on Applied computing**

Publisher: ACM Press

Full text available:  pdf(618.92 KB) Additional Information: [full citation](#), [references](#), [index terms](#)

Keywords: RAID, clusters of workstations, false sharing, storage system architecture

5 The HP AutoRAID hierarchical storage system

 John Wilkes, Richard Golding, Carl Staelin, Tim Sullivan

February 1996 **ACM Transactions on Computer Systems (TOCS)**, Volume 14 Issue 1

Publisher: ACM Press

Full text available:  pdf(1.82 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Configuring redundant disk arrays is a black art. To configure an array properly, a system administrator must understand the details of both the array and the workload it will support. Incorrect understanding of either, or changes in the workload over time, can lead to poor performance. We present a solution to this problem: a two-level storage hierarchy implemented inside a single disk-array controller. In the upper level of this hierarchy, two copies of active data are stored to provide f ...

Keywords: RAID, disk array, storage hierarchy

6 X-RAY: A Non-Invasive Exclusive Caching Mechanism for RAIDs

 Lakshmi N. Bairavasundaram, Muthian Sivathanu, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau

March 2004 **ACM SIGARCH Computer Architecture News , Proceedings of the 31st annual international symposium on Computer architecture ISCA '04**, Volume 32 Issue 2

Publisher: IEEE Computer Society, ACM Press

Full text available:  pdf(250.59 KB) Additional Information: [full citation](#), [abstract](#), [citations](#)

RAID storage arrays often possess gigabytes of RAM for caching disk blocks. Currently, most RAID systems use LRU or LRU-like policies to manage these caches. Since these array caches do not recognize the presence of file system buffer caches, they redundantly retain many of the same blocks as those cached by the file system, thereby wasting precious cache space. In this paper, we introduce X-RAY, an exclusive RAID array caching mechanism. X-RAY achieves a high degree of (but not perfect) exclusivity through ...

7 Virtual memory and backing storage management in multiprocessor operating systems using object-oriented design techniques



V. F. Russo, R. H. Campbell.

September 1989 **ACM SIGPLAN Notices , Conference proceedings on Object-oriented programming systems, languages and applications OOPSLA '89,**
Volume 24 Issue 10

Publisher: ACM Press

Full text available: pdf(1.19 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The Choices operating system architecture [3, 4, 15] uses class hierarchies and object-oriented programming to facilitate the construction of customized operating systems for shared memory and networked multiprocessors. The software is being used in the Tapestry Parallel Computing Laboratory at the University of Illinois to study the performance of algorithms, mechanisms, and policies for parallel systems. This paper describes the architectural design and class hierarchy of ...

8 Computing curricula 2001



September 2001 **Journal on Educational Resources in Computing (JERIC)**

Publisher: ACM Press

Full text available: pdf(613.63 KB)

html(2.78 KB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

9 Disk cache—miss ratio analysis and design considerations



Alan J. Smith

August 1985 **ACM Transactions on Computer Systems (TOCS)**, Volume 3 Issue 3

Publisher: ACM Press

Full text available: pdf(3.13 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

The current trend of computer system technology is toward CPUs with rapidly increasing processing power and toward disk drives of rapidly increasing density, but with disk performance increasing very slowly if at all. The implication of these trends is that at some point the processing power of computer systems will be limited by the throughput of the input/output (I/O) system. A solution to this problem, which is described and evaluated in this paper, is disk cache

10 The HP AutoRAID hierarchical storage system



J. Wilkes, R. Golding, C. Staelin, T. Sullivan

December 1995 **ACM SIGOPS Operating Systems Review , Proceedings of the fifteenth ACM symposium on Operating systems principles SOSP '95**, Volume 29
Issue 5

Publisher: ACM Press

Full text available: pdf(1.60 MB)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

11 The Zebra striped network file system



John H. Hartman, John K. Ousterhout

August 1995 **ACM Transactions on Computer Systems (TOCS)**, Volume 13 Issue 3

Publisher: ACM Press

Full text available: pdf(2.76 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Zebra is a network file system that increases throughput by striping the file data across multiple servers. Rather than striping each file separately, Zebra forms all the new data

from each client into a single stream, which it then stripes using an approach similar to a log-structured file system. This provides high performance for writes of small files as well as for reads and writes of large files. Zebra also writes parity information in each stripe in the style of RAID disk arrays; this ...

Keywords: RAID, log-based striping, log-structured file system, parity computation

12 The automatic improvement of locality in storage systems

 Windsor W. Hsu, Alan Jay Smith, Honesty C. Young

November 2005 **ACM Transactions on Computer Systems (TOCS)**, Volume 23 Issue 4

Publisher: ACM Press

Full text available:  pdf(2.58 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Disk I/O is increasingly the performance bottleneck in computer systems despite rapidly increasing disk data transfer rates. In this article, we propose Automatic Locality-Improving Storage (ALIS), an introspective storage system that automatically reorganizes selected disk blocks based on the dynamic reference stream to increase effective storage performance. ALIS is based on the observations that sequential data fetch is far more efficient than random access, that improving seek distances prod ...

Keywords: Data layout optimization, block layout, data reorganization, data restructuring, defragmentation, disk technology trends, locality improvement, prefetching

13 LH*⊂RS---a highly-available scalable distributed data structure

 Witold Litwin, Rim Moussa, Thomas Schwarz

September 2005 **ACM Transactions on Database Systems (TODS)**, Volume 30 Issue 3

Publisher: ACM Press

Full text available:  pdf(774.32 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

LH*RS is a high-availability scalable distributed data structure (SDDS). An LH*RS file is hash partitioned over the distributed RAM of a multicomputer, for example, a network of PCs, and supports the unavailability of any $k \geq 1$ of its server nodes. The value of k transparently grows with the file to offset the reliability decline. Only the number of the storage nodes potentially limits the file growth. The high-availability management uses a novel ...

Keywords: P2P, Scalable distributed data structure, grid computing, high-availability, linear hashing, physical database design

14 Compiler-based I/O prefetching for out-of-core applications

 Angela Demke Brown, Todd C. Mowry, Orran Krieger

May 2001 **ACM Transactions on Computer Systems (TOCS)**, Volume 19 Issue 2

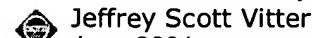
Publisher: ACM Press

Full text available:  pdf(499.03 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Current operating systems offer poor performance when a numeric application's working set does not fit in main memory. As a result, programmers who wish to solve "out-of-core" problems efficiently are typically faced with the onerous task of rewriting an application to use explicit I/O operations (e.g., read/write). In this paper, we propose and evaluate a fully automatic technique which liberates the programmer from this task, provides high performance, and requires only minima ...

Keywords: compiler optimization, prefetching, virtual memory

15 [External memory algorithms and data structures: dealing with massive data](#)



Jeffrey Scott Vitter

June 2001 **ACM Computing Surveys (CSUR)**, Volume 33 Issue 2

Publisher: ACM Press

Full text available: [pdf\(828.46 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Data sets in large applications are often too massive to fit completely inside the computers internal memory. The resulting input/output communication (or I/O) between fast internal memory and slower external memory (such as disks) can be a major performance bottleneck. In this article we survey the state of the art in the design and analysis of external memory (or EM) algorithms and data structures, where the goal is to exploit locality in order to reduce the I/O costs. We consider a varie ...

Keywords: B-tree, I/O, batched, block, disk, dynamic, extendible hashing, external memory, hierarchical memory, multidimensional access methods, multilevel memory, online, out-of-core, secondary storage, sorting

16 [I/O reference behavior of production database workloads and the TPC benchmarks— an analysis at the logical level](#)



Windsor W. Hsu, Alan Jay Smith, Honesty C. Young

March 2001 **ACM Transactions on Database Systems (TODS)**, Volume 26 Issue 1

Publisher: ACM Press

Full text available: [pdf\(5.42 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

As improvements in processor performance continue to far outpace improvements in storage performance, I/O is increasingly the bottleneck in computer systems, especially in large database systems that manage huge amoungs of data. The key to achieving good I/O performance is to thoroughly understand its characteristics. In this article we present a comprehensive analysis of the logical I/O reference behavior of the peak productiondatabase workloads from ten of the world's largest corporatio ...

Keywords: I/O, TPC benchmarks, caching, locality, prefetching, production database workloads, reference behavior, sequentiality, workload characterization

17 [Application performance and flexibility on exokernel systems](#)



M. Frans Kaashoek, Dawson R. Engler, Gregory R. Ganger, Hector M. Briceño, Russell Hunt,

David Mazières, Thomas Pinckney, Robert Grimm, John Jannotti, Kenneth Mackenzie

October 1997 **ACM SIGOPS Operating Systems Review , Proceedings of the sixteenth ACM symposium on Operating systems principles SOSP '97**, Volume 31 Issue 5

Publisher: ACM Press

Full text available: [pdf\(2.39 MB\)](#)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

18 [The Zebra striped network file system](#)



John H. Hartman, John K. Ousterhout

December 1993 **ACM SIGOPS Operating Systems Review , Proceedings of the fourteenth ACM symposium on Operating systems principles SOSP '93**, Volume 27 Issue 5

Publisher: ACM Press

Full text available:  pdf(1.93 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Zebra is a network file system that increases throughput by striping file data across multiple servers. Rather than striping each file separately, Zebra forms all the new data from each client into a single stream, which it then stripes using an approach similar to a log-structured file system. This provides high performance for writes of small files as well as for reads and writes of large files. Zebra also writes parity information in each stripe in the style of RAID disk arrays; this increase ...

19 Asynchronous scheduling of redundant disk arrays 

 Peter Sanders

July 2000 **Proceedings of the twelfth annual ACM symposium on Parallel algorithms and architectures**

Publisher: ACM Press

Full text available:  pdf(161.35 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Random redundant allocation of data to parallel disk arrays can be exploited to achieve low access delays. New algorithms are proposed which improve the previously known shortest queue algorithm by systematically exploiting that scheduling decisions can be deferred until a block access is actually started on a disk. These algorithms are also generalized for coding schemes with low redundancy. Using extensive experiments, practically important quantities are measured which have so far eluded ...

20 General storage protection techniques: Ensuring data integrity in storage: techniques and applications 

 Gopalan Sivathanu, Charles P. Wright, Erez Zadok

November 2005 **Proceedings of the 2005 ACM workshop on Storage security and survivability StorageSS '05**

Publisher: ACM Press

Full text available:  pdf(217.83 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Data integrity is a fundamental aspect of storage security and reliability. With the advent of network storage and new technology trends that result in new failure modes for storage, interesting challenges arise in ensuring data integrity. In this paper, we discuss the causes of integrity violations in storage and present a survey of integrity assurance techniques that exist today. We describe several interesting applications of storage integrity checking, apart from security, and discuss the im ...

Keywords: file systems, intrusion detection, storage integrity

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)



[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

Search: The ACM Digital Library The Guide

SEARCH

THE ACM DIGITAL LIBRARY

[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Petal: distributed virtual disks

Full text [Pdf \(1.10 MB\)](#)

Source [Architectural Support for Programming Languages and Operating Systems archive](#)

[Proceedings of the seventh international conference on Architectural support for programming languages and operating systems table of contents](#)

Cambridge, Massachusetts, United States

Pages: 84 - 92

Year of Publication: 1996

ISBN:0-89791-767-7

[Also published in ...](#)

Authors [Edward K. Lee](#) Systems Research Center, Digital Equipment Corporation, 130 Lytton Ave, Palo Alto, CA
[Chandramohan A. Thekkath](#) Systems Research Center, Digital Equipment Corporation, 130 Lytton Ave, Palo Alto, CA

Sponsors [SIGPLAN](#): ACM Special Interest Group on Programming Languages

[SIGOPS](#): ACM Special Interest Group on Operating Systems

[IEEE-CS](#) : Computer Society

[SIGARCH](#): ACM Special Interest Group on Computer Architecture

Publisher ACM Press New York, NY, USA

[Sign in](#)
[Web](#) [Images](#) [Groups](#) [News](#) [Froogle](#) [Local](#) [more »](#)

1999 disk pack algorithm

[Advanced Search](#)
[Preferences](#)
WebResults 1 - 10 of about 903,000 for **1999 disk pack algorithm**. (0.58 seconds)**A Disk-Packing Algorithm for an (ResearchIndex)**

42.7%: A Disk-Packing Algorithm for an Origami Magic Trick - Bern, Demaine, ... 0.3:
 Ununfoldable Polyhedra - Bern, Demaine, Eppstein, Kuo (1999) (Correct) ...
citeseer.ist.psu.edu/652666.html - 17k - [Cached](#) - [Similar pages](#)

Citations: A disk-packing algorithm for an origami magic trick ...

A disk-packing algorithm for an origami magic trick. ... Eppstein, Erickson (1999) Self-citation (Eppstein) (Correct) from which it was derived, ...
citeseer.ist.psu.edu/context/331766/383052 - 15k - [Cached](#) - [Similar pages](#)

Circle Packing -- From MathWorld

Collins, CR and Stephenson, K. "A Circle Packing Algorithm." Preprint. ... Stephenson, K.
 "Circle Packing Bibliography as of April 1999. ...
mathworld.wolfram.com/CirclePacking.html - 53k - [Cached](#) - [Similar pages](#)

Folding -- From MathWorld

... D.; and Hayes, B. "A Disk-Packing Algorithm for an Origami Magic Trick. ... 891-892,
1999. Gallivan, BC "How to Fold Paper in Half Twelve Times: An ...
mathworld.wolfram.com/Folding.html - 35k - [Cached](#) - [Similar pages](#)

David Eppstein - Publications

J. Algorithms 30:302-322, **1999** (special issue for SODA 1997). ... A disk-packing
 algorithm for an origami magic trick. M. Bern, E. Demaine, D. Eppstein, ...
www.ics.uci.edu/~eppstein/pubs/1999.html - 26k - [Cached](#) - [Similar pages](#)

The Fold-and-Cut Problem (Erik Demaine)

I call them the straight-skeleton and disk-packing solutions. ... Freie Universität Berlin
 (December **1999**), Symposium on Discrete Algorithms (January **1999**), ...
theory.ics.mit.edu/~edemaine/foldcut/ - 14k - [Cached](#) - [Similar pages](#)

Citations: On multidimensional packing problems - Chekuri, Khanna ...

C. Chekuri and S. Khanna. On multidimensional packing problems. In ACM Symp. on
 Discrete Algorithms, pages 185–194, **1999**.
citeseer.ifi.unizh.ch/context/949213/463513 - 24k - [Cached](#) - [Similar pages](#)

1999 USENIX Annual Technical Conference, June 6-11, **1999, Monterey ...**

The common bottleneck of large caching proxy servers is disk I/O. In this ... We use a
 packing algorithm, shown in Algorithm 3 that for the given traces ...
www.usenix.org/events/usenix99/full_papers/maltzahn/maltzahn_html/index.html - 81k -
[Cached](#) - [Similar pages](#)

Generating Conformal Flat Maps of the Cortical Surface via Circle ...

... representation of a cortical surface is topologically equivalent to a disk ... a new
 algorithm which computes the circle packing of a simply connected ...
www.math.fsu.edu/~mhurdal/posters/hbm99math.html - 30k - [Cached](#) - [Similar pages](#)

[PDF] arXiv:cs.CG/9908006 v2 13 Aug **1999**

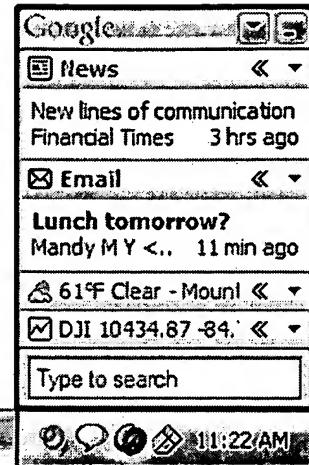
File Format: PDF/Adobe Acrobat - [View as HTML](#)
 CG/9908006 v2 13 Aug **1999**. Computational Geometry Column 36 ... A disk-packing

algorithm for an origami. magic trick. In Proc. Internat. Conf. ...
arxiv.org/pdf/cs.CG/9908006 - [Similar pages](#)

Try your search again on [Google Book Search](#)

Gooooooooooooogle ►
Result Page: 1 2 3 4 5 6 7 8 9 10 [Next](#)

Info when you want it, right on your desktop
Free! [Download Google Desktop](#)



[Search within results](#) | [Language Tools](#) | [Search Tips](#) | [Dissatisfied? Help us improve](#)

[Google Home](#) - [Advertising Programs](#) - [Business Solutions](#) - [About Google](#)

©2006 Google

[Sign in](#)
[Web](#) [Images](#) [Groups](#) [News](#) [Froogle](#) [Local](#) [more »](#)

1999 epoch disk pack algorithm

[Advanced Search](#)
[Preferences](#)
WebResults 21 - 30 of about 38,700 for **1999 epoch disk pack algorithm**. (0.65 seconds)

[\[PPT\] Lecture 23: Remote Procedure Call](#)

File Format: Microsoft Powerpoint 97 - [View as HTML](#)

Epoch # stored on **disk**, Put in every message; epoch # incremented on crash and/or when ... Code for client to **pack** message, send it off, wait for result, ...

www-inst.eecs.berkeley.edu/~cs162/sp06/Lectures/lec23-rpc.ppt - [Similar pages](#)

[\[PDF\] Page 1 CS162 Operating Systems and Systems Programming Lecture 23 ...](#)

File Format: PDF/Adobe Acrobat - [View as HTML](#)

Sender uses an adaptive **algorithm** to decide size of N. » Goal: fill network between sender and receiver ... Epoch # stored on **disk**, Put in every message ...

www-inst.eecs.berkeley.edu/~cs162/sp06/Lectures/lec23-rpc.pdf - [Similar pages](#)

[\[PDF\] Information and Control in Gray-Box Systems](#)

File Format: PDF/Adobe Acrobat - [View as HTML](#)

from **disk**, due to the LRU-like page replacement **algorithm**. The gray-box scan is able to ...

Disk Characteristics. Berkeley TR CSD-99-1063, 1999. ...

sosp.org/2001/papers/arpacidusseau.pdf - [Similar pages](#)

[Epoch Audio Inc.](#)

Laser Disc: Now-defunct 12-inch **disc** format with excellent analog, FM-recorded video image, ... Epoch Audio Inc.(China) All rights reserved ©1999-2006.

www.epochaudio.com/support/pro_reference_links.htm - 61k - [Cached](#) - [Similar pages](#)

[\[PDF\] Software Release Bulletin: VOS Release 14.3.0](#)

File Format: PDF/Adobe Acrobat - [View as HTML](#)

7200 RPM SCSI **disk** drives— As of September 1, 1999 the D803, D705, D706, ... 13:03:08

Renaming **disk pack** m1_d01.0-3.pri to m1_d05.0-3.pri. ...

ftp.stratus.com/vos/srbs/r914-03.pdf - [Similar pages](#)

[\[PDF\] System Support for Background Replication](#)

File Format: PDF/Adobe Acrobat - [View as HTML](#)

This **algorithm** is designed to support massive replication. of data and services, where hardware (eg, bandwidth, **disk**, space, and processor cycles) is ...

www.cs.cornell.edu/People/egs/syslunch-fall02/bgrep.pdf - [Similar pages](#)

[\[PDF\] Managing Prefetch Memory for Data-Intensive Online Servers](#)

File Format: PDF/Adobe Acrobat - [View as HTML](#)

Technical Report. HPL-1999-35, HP Laboratories Palo Alto, 1999. [3] Ask Jeeves Search.

<http://www.ask.com>. [4] LA Belady. A Study of Replacement **Algorithms** ...

www.cs.rochester.edu/~kshen/papers/fast2005-li.pdf - [Similar pages](#)

[SwRI/BEO 1997 Research Highlights](#)

During 1998 and 1999, SwRI scientists and engineers will build two ALICE prototypes, ...

This miniaturized, high-tech spectrometer **packs** into a 3 kg, ...

www.boulder.swri.edu/research.html - 13k - [Cached](#) - [Similar pages](#)

[\[PDF\] jcn99714 363..375](#)

File Format: PDF/Adobe Acrobat - [View as HTML](#)

for the **disc** stimulus and the control **epoch** (p < .01., uncorrected). ... Glover, GH (1999).

Simple analytic spiral K-space **algorithm**. ...
www.brown.edu/Research/LCE/fMRIJC/Mendola_JCN_2006.pdf - [Similar pages](#)

[ps] [NPS: A Non-interfering Deployable Web Prefetching System \Lambda](#)mbda ...

File Format: Adobe PostScript - [View as HTML](#)

In each **epoch**, the monitor collects a response time sample for a ... A Trace-Driven Comparison of **Algorithms** for Parallel Prefetching and Caching. ...
www.cs.utexas.edu/users/rkoku/RESEARCH/NPS.ps - [Similar pages](#)

◀ Goooooooooooooogle ▶

Result Page: [Previous](#) [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [11](#) [12](#) [Next](#)

[Search within results](#) | [Language Tools](#) | [Search Tips](#)

[Google Home](#) - [Advertising Programs](#) - [Business Solutions](#) - [About Google](#)

©2006 Google

[Sign in](#)[Web](#) [Images](#) [Groups](#) [News](#) [Froogle](#) [Local](#) [more »](#)

1999 raid epoch identifiers

[Advanced Search](#)
[Preferences](#)**Web**Results 1 - 10 of about 10,500 for **1999 raid epoch identifiers**. (0.49 seconds)**[PDF] NCC 98 Initial Release Operational Readiness Review (ORR) January ...**File Format: PDF/Adobe Acrobat - [View as HTML](#)

overlapping epoch times that are received with less than 30 second ... another Support Identifier (SUPIDEN) was to follow but one did ...

ncc.gsfc.nasa.gov/hcc98/trans/orr-pres.pdf - [Similar pages](#)**New Debian i386 Packages**

... urgency=low * changed epoch. dbuild - Tool for building Debian binary packages from ...

Sendmail+UUCP Sendmail-Address-Rewrite Software-RAID VAIO+Linux ...

lists.debian.org/debian-changes/1999/03/msg00052.html - 60k - [Cached](#) - [Similar pages](#)**New Debian alpha Packages**

... urgency=low * changed epoch. dbuild - Tool for building Debian binary packages ... In man-db -69a: Sun, 7 Mar 1999 19:32:25 +0200 * Applied (manually) ...

lists.debian.org/debian-changes/1999/03/msg00051.html - 44k - [Cached](#) - [Similar pages](#)[[More results from lists.debian.org](#)]**[doc] Version 0 System Upgrade**File Format: Microsoft Word 2000 - [View as HTML](#)In February of 1999, the NOFS was replaced by a Sun E4500 with an additional 180 gigabyte Redundant Array of Inexpensive Disks (**RAID**). ...www.space.qinetiq.com/ceos/docs/ USGS%20Browse%20Study%20August%202001.doc - [Similar pages](#)**[PDF] Manageability, Availability and Performance in Porcupine: A Highly ...**File Format: PDF/Adobe Acrobat - [View as HTML](#)

Networks 1999], although nothing in the system prevents their use. ... coordinator broadcasts the new membership and epoch ID to all nodes. ...

www.cs.washington.edu/homes/levy/porcupine.pdf - [Similar pages](#)**Thu May 10 17:24:54 PDT 2001 * Release Orca version 0.27b1. Thu ...**Wed Oct 20 17:37:01 PDT 1999 * Fix that annoying warning from pod2html when orca.html ... Now all classes use File **IDentifiers** (FIDs) instead of filenames. ...www.orcaware.com/orca/pub/CHANGES-0.27b1 - 49k - [Cached](#) - [Similar pages](#)**NEW IN ORCA 0.27 ===== Released November 14, 2002. 1 ...**The subroutine should return the Unix epoch time. ... 11) Use a new integer file **identifier** (FID) instead of the filename to refer to in internal objects. ...www.orcaware.com/svn/repos/orca/trunk/NEWS - 35k - [Cached](#) - [Similar pages](#)**[PDF] Fault-tolerant replication management in large-scale distributed ...**File Format: PDF/Adobe Acrobat - [View as HTML](#)

the epoch number at each change in the layout, and ensures. that the manager and all non-failed ... building distributed RAID-5 disk arrays. The data access ...

www.hpl.hp.com/personal/John_Wilkes/papers/SRDS99.pdf - [Similar pages](#)**[PDF] Records Management, Preservation, Indexing, and Access Plan ...**File Format: PDF/Adobe Acrobat - [View as HTML](#)structure of the Internet and intranet servers **RAID** (Redundant Array of Inexpensive Disks).

arrays. In Epoch 2, each fascicle (for example, a DVD) contains ...
navigatela.lacity.org/.../vault_record/ Records%20Management%20Plan%
20v010.pdf&category=vault_record - [Similar pages](#)

[ps] [Manageability, Availability and Performance in Porcupine: A Highly ...](#)

File Format: Adobe PostScript - [View as HTML](#)

to use a large monolithic server with reliable storage (eg, RAID [Chen et al. ... Each entry in
the user map is associated with an **epoch** ID that shows when ...

[www.ysaito.com/tocs.ps](#) - [Similar pages](#)

Try your search again on [Google Book Search](#)

Gooooooooogle ►

Result Page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [Next](#)

Free! Speed up the web. [Download the Google Web Accelerator.](#)

[Search within results](#) | [Language Tools](#) | [Search Tips](#) | [Dissatisfied? Help us improve](#)

[Google Home](#) - [Advertising Programs](#) - [Business Solutions](#) - [About Google](#)

©2006 Google

[Sign in](#)
[Web](#) [Images](#) [Groups](#) [News](#) [Froogle](#) [Local](#) [more »](#)

1999 disk epoch raid cluster logical volume

[Advanced Search](#)
[Preferences](#)
WebResults 1 - 10 of about 9,310 for **1999 disk epoch raid cluster logical volume** . (0.67 seconds)**Scholarly articles for 1999 disk epoch raid cluster logical volume**[Manageability, Availability, and Performance in ...](#) - by Saito - 104 citations[The Lustre Storage Architecture](#) - by Braam - 31 citations[Storage-Aware Caching: Revisiting Caching for ...](#) - by Forney - 13 citations[\[PDF\] Manageability, Availability and Performance in Porcupine: A Highly ...](#)File Format: PDF/Adobe Acrobat - [View as HTML](#)

and manageability at high **volume** are demanded. For example, Usenet news, community ... We then configured a two node **cluster**, each with one IDE **disk** and ...
www.cs.washington.edu/homes/levy/porcupine.pdf - [Similar pages](#)

[Citations: The Logical Disk: A New Approach to Improving File ...](#)

Related block level storage systems include **RAID** II [7] TickerTAIP [5] **Logical Disk** [8], Loge [10] Mime [6] AutoRAID [19] and Swift [4] Some of these ...
citeseer.ifi.unizh.ch/context/135725/0 - 44k - [Cached](#) - [Similar pages](#)

[\[PDF\] Differentiated Object Placement and Location for Self-Organizing ...](#)

File Format: PDF/Adobe Acrobat

a **RAID** system. **Volume** formation information, called vol- ... When a provider detects a **disk** failure, it also increments the **epoch** num- ...
www.cs.ucsb.edu/research/tech_reports/reports/2002-32.pdf - [Similar pages](#)

[The HP AutoRAID hierarchical storage system](#)Striping in a **RAID** level-5 **disk** array. Tech. Pep. CSE-TR- 181-93, ... 20 **EPOCH**

SYSTEMS. 1988. Mass storage: Server puts optical discs on line for ...

portal.acm.org/citation.cfm?id=225539&[dl=ACM&coll=GUIDE&CFID=15151515&CFTOKEN=6184618](#) - [Similar pages](#)[\[PDF\] Manageability, Availability, and Performance in Porcupine: A ...](#)File Format: PDF/Adobe Acrobat - [View as HTML](#)

configured a two-node **cluster**, each with one IDE **disk** and two SCSI disks. ... 1998; Sun Microsystems 1999]). These **clusters** are often designed specifi- ...
www.cs.nyu.edu/rgrimm/teaching/readings/porcupine.pdf - [Similar pages](#)

[The Big Brother FTP Archive Index](#)... to check for stale **logical volumes** and volumegroups for unavailable disks ...**1999/10/20, 1418, ods.sh, A script to monitor Suns Online Disk Suite RAID.** ...www.deadcat.net.au/cgi-bin/BB.pl - 249k - Apr 15, 2006 - [Cached](#) - [Similar pages](#)[\[PDF\] Track-Based Disk Logging](#)

File Format: PDF/Adobe Acrobat

Trail log **disk** when the Trail driver boots up. The **epoch**. field is incremented each time when ... **clusters** **disk** access activities that could have been more ...

doi.ieeecomputersociety.org/10.1109/DSN.2002.1028928 - [Similar pages](#)[\[PS\] SEMANTICALLY-SMART DISK SYSTEMS by Muthian Sivathanu BE Computer ...](#)File Format: Adobe PostScript - [View as HTML](#)

For example, even 2 failures in a 10 **disk** **RAID** system would result in ... In such a scenario, D-GRAID can be run at the **logical volume manager** level, ...

www.cs.wisc.edu/adsl/Publications/phd-muthian.ps - [Similar pages](#)

Linux Weekly News

Ciprico Inc. (MINNEAPOLIS) announced that its 6500 Series **RAID disk array** is now ...

Linux **Logical Volume Management** HOWTO, 0.0.1, A HOWTO describing Linux ...

lwn.net/2000/0406/bigpage.php3 - 187k - [Cached](#) - [Similar pages](#)

[PDF] [IBM E server Certification Study Guide: Cluster 160000 Managed by PSSP](#)

File Format: PDF/Adobe Acrobat

RAID disk storage. Customers with storage needs that are ... Aggregates cannot be larger than the **logical volume** in which they. are contained. ...

www.redbooks.ibm.com/redbooks/pdfs/sg247013.pdf - [Similar pages](#)

Try your search again on [Google Book Search](#)

Gooooooooogle ►

Result Page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [Next](#)

Free! Speed up the web. [Download the Google Web Accelerator](#).

[Search within results](#) | [Language Tools](#) | [Search Tips](#) | [Dissatisfied? Help us improve](#)

[Google Home](#) - [Advertising Programs](#) - [Business Solutions](#) - [About Google](#)

©2006 Google

[Sign in](#)
[Web](#) [Images](#) [Groups](#) [News](#) [Froogle](#) [Local](#) [more »](#)

1999 backup epoch

[Advanced Search](#)
[Preferences](#)
WebResults 1 - 10 of about 234,000 for **1999 backup epoch** . (0.33 seconds)**Storage Mountain: AMANDA Book Excerpt 23**

This is an excerpt from the book Unix **Backup & Recovery**. ... print &ctime (\$_) ; } exit (0); \$ epoch.pl 916549924 Sun Jan 17 0:12:04 US/East-Indiana **1999** ...
www.backupcentral.com/amanda-23.html - 20k - Cached - Similar pages

1999: Summary: Backup - ufsdump question

The idea is that ufsdump will only **backup** more than one file / directory if ... to back the two filesystems on the first tape drive, **epoch** level, eject ...
www.sunmanagers.org/archives/1999/1822.html - 6k - Cached - Similar pages

Palmdump

Most Unix and POSIX-compatible programs use a different **epoch** or date and time ... **1999**
 Modification time: Palm time: Mon Jul 12 19:31:39 **1999** Backup time: ...
www.fourmilab.ch/palm/palmdump/ - 13k - Cached - Similar pages

docs.sun.com: System Administration Guide, Volume 1

... 13 10:46:09 **1999** DUMP: Date of last level 0 dump: the **epoch** DUMP: Dumping ... The following example shows an incremental **backup** of the root (/) file ...
docs.sun.com/app/docs/doc/805-7228/6j6q7uf1d?a=view - 29k - Cached - Similar pages

docs.sun.com: System Administration Guide, Volume 1

Preparing to **back up** file systems begins with planning, ... level 0 dump: Tue Jul 13 10:46:09 **1999** DUMP: Date of last level 0 dump: the **epoch** DUMP: Dumping ...
docs.sun.com/app/docs/doc/805-7228/6j6q7uf1a?a=view - 34k - Cached - Similar pages

BackupFAQs < CSDocs < TWiki

Incremental and **epoch** backups (see What Are Backups?) are made from the **backup** volume. This means that a file from an AFS volume can only be restored to ...
www.cs.wisc.edu/twiki/bin/view/CSDocs/BackupFAQs - 15k - Cached - Similar pages

CDDIS | Bulletin Volume 15, No. 3

Epoch: March 15, 1992, TLRS-1 range bias is +11 mm range bias ... Since the December 1999 issue of the CDDIS Bulletin, the CDDIS has archived daily and ...
cddis.nasa.gov/bulletin_v15n3.html - 38k - Cached - Similar pages

Free Epoch Convertor

Free **Epoch** Convertor, GMT converter, convert seconds to date, birthdate, ... The **Epoch** started January 1st, 1970 at 00:00:01. The **Epoch** is the number of ...
www.perlservices.net/en/programs/epoch_converter/epoch_converter.html - 20k - Cached - Similar pages

[PDF] GEOULINKSUBSYSTEM(GUS) CLOCKSTEERINGALGORITHMS PERFORMANCE ...

File Format: PDF/Adobe Acrobat - [View as HTML](#)

Time (days) past UTC 00:00:00, 1/06/**1999**. Figure 4: Primary GUS Processing (GP)
 Clock Steering for AOR-W (Clarksburg, MD). **Backup** GP Clock Steering ...
tycho.usno.navy.mil/ptti/ptti99/PTTI_1999_173.PDF - Similar pages

System and method for restoring a file system from backups in the ...

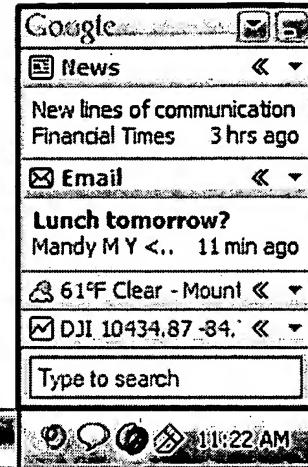
5974563, Oct., **1999**, Beeler, Jr. 6073128, Jun., 2000, Pongracz et al. ... Deletion events for

the files that have survived a **back-up epoch** are recorded in a ...
www.freepatentsonline.com/6938056.html - 38k - [Cached](#) - [Similar pages](#)

Try your search again on [Google Book Search](#)

Gooooooooooooogle ►
Result Page: 1 2 3 4 5 6 7 8 9 10 [Next](#)

Info when you want it, right on your desktop
Free! [Download Google Desktop](#)



[Search within results](#) | [Language Tools](#) | [Search Tips](#) | [Dissatisfied? Help us improve](#)

[Google Home](#) - [Advertising Programs](#) - [Business Solutions](#) - [About Google](#)

©2006 Google

[Sign in](#)
[Web](#) [Images](#) [Groups](#) [News](#) [Froogle](#) [Local](#) [more »](#)

1999 raid disk backup epoch

[Search](#)
[Advanced Search](#)
[Preferences](#)
WebResults 1 - 10 of about 21,600 for **1999 raid disk backup epoch** . (0.64 seconds)**backups storage data on GlobalSpec**

RAID Disk Arrays (56 companies) RAID disk arrays are collections of ...

ITnetcentral - Research - IT Library: **Disk Backups**, Storage Management Services. ...

data-storage-systems.globalspec.com/ Industrial-Directory/backups_storage_data - 79k - [Cached](#) - [Similar pages](#)

Sponsored Links**Raid Backup**

Arctor **Backup & Version Management**
 Fast. Secure. Simple. - Free Trial!
www.byteplant.com/products/arctor

data storages external on GlobalSpec

RAID Disk Arrays (56 companies) RAID disk arrays are collections of storage disks

... All-In-One RAID 5 terabyte storage Auto **backup**, networked and secure. ...

data-storage-systems.globalspec.com/ Industrial-Directory/data_storages_external - 80k - [Cached](#) - [Similar pages](#)

[PS] Astronomical Data Analysis Software and Systems VIII ASP ...

File Format: Adobe PostScript - [View as HTML](#)

269 c fl Copyright 1999 Astronomical Society of the Pacific. ... **disk** farm lots of large **disks**, **backup** on DLT \$64/GB yes RAID farm lots of (level-5) RAID ...

www.adass.org/adass/proceedings/adass98/reprints/pirenneb.ps.gz - [Similar pages](#)

The Prospects of DVD-R for Storing Astronomical Archive Data

Magnetic **disk** systems are RAID (Redundant Array of Inexpensive Disks) systems: their cost is increased to account for a necessary **backup** on say tapes or for ...

www.adass.org/adass/proceedings/adass98/pirenneb/ - 17k - [Cached](#) - [Similar pages](#)

[PDF] Resume for Jordan Schwartz

File Format: PDF/Adobe Acrobat - [View as HTML](#)

with specialization in **backup** and recovery, storage, database, and ... system **disks** and Winchester RAID Array. Detailed build notes were ...

home.adelphia.net/~jordan247/resume/resume.pdf - [Similar pages](#)

[PDF] Reducing Energy Consumption of Disk Storage Using Power-Aware ...

File Format: PDF/Adobe Acrobat - [View as HTML](#)

ing the previous epoch for each disk. (1) If the percentage of ... Grunwald [8] proposed a **disk backup** organization called MAID that uses power management ... opera.cs.uiuc.edu/paper/HPCA04.pdf - [Similar pages](#)

[PDF] €¡¤£!¥"§ © ¡¤ © ! "§# ¥\$©&%' ()¡1032 4 "56 798 ...

File Format: PDF/Adobe Acrobat - [View as HTML](#)

Is it acceptable to lose any data if a **disk** failure ... In the epoch algorithm [4] each **backup** server ... separate RAID devices [5]. These systems are ...

www.it.nuigalway.ie/Publications/TR/abstracts/NUIG-IT-190900.pdf - [Similar pages](#)

Power-Aware Storage Cache Management

In each epoch, we keep track of the interval length between two consecutive ...

Colarelli and Grunwald [8] proposed a **disk backup** organization called MAID ...

doi.ieeecomputersociety.org/10.1109/TC.2005.82 - [Similar pages](#)

Citations: The Logical Disk: A New Approach to Improving File ...

None of these supported multiple storage levels, and none was targeted at **RAID** arrays. Work on an Extended Function Controller at HP's **disk** divisions in the ...
citeseer.ist.psu.edu/context/135725/0 - 37k - Cached - Similar pages

Epinions.com - Member Center:

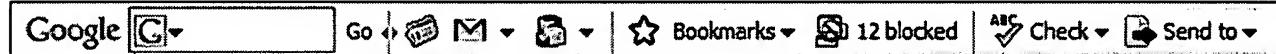
Clarion CX ATA **DISK ARRAY EXPAN** HOLDS 15 ATA **DISK DR** (cx-ata-dae) · Clarion CX
FC **DISK ARRAY** ... **Epoch** 324GB 9bay Rkmt **RAID** VHDCI U160 Dual (e9ir324gbdb) ...
www.epinions.com/browse/new_products/ new_prod_cat_~158963073/category_~Network+Storage - 87k -
Cached - Similar pages

Try your search again on [Google Book Search](#)

Gooooooooooooogle ►

Result Page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [Next](#)

Free! Get the Google Toolbar. [Download Now](#) - [About Toolbar](#)



[Search within results](#) | [Language Tools](#) | [Search Tips](#) | [Dissatisfied? Help us improve](#)

[Google Home](#) - [Advertising Programs](#) - [Business Solutions](#) - [About Google](#)

©2006 Google

We are experiencing some difficulties with our database and, consequently, some document links may be incorrect. Please bear with us while we fix the problem.

31 citations found. Retrieving documents...

W. Jonge, M. F. Kaashoek, and W. C. Hsieh. *The Logical Disk: A New Approach to Improving File Systems*. In Proceedings of the Fourteenth Symposium on Operating Systems Principles, 1993.

CiteSeer [Home/Search](#) Document Not in Database [Summary](#) [Related Articles](#) [Check](#)

This paper is cited in the following contexts:

First 50 documents

[An Age-Threshold Algorithm for Garbage Collection in..](#) - Menon, Stockmeyer (1998) (1 citation) (Correct)

....the file system. **Because** LSA is implemented in an outboard controller which has no understanding of files, it is more appropriate to think of LSA as a log structured track manager rather than a log structured file system. **In this respect, LSA has some similarities to Loge [4] and to Logical Disk [3], both of which are implemented below the file system.** In LSA, updated data is written into new disk locations instead of being written in place. Large amounts of updated data are collected in controller memory and written together to a contiguous segment on the disks. **Parity** on this data is also

W. de Jonge, M. F. Kaashoek, and W. C. Hsieh, *The logical disk: a new approach to improving file systems*, Proc. 14th ACM Symposium on Operating System Principles, Dec. 1993, Asheville, NC, pp. 15–28.

[FT-NFS: an Efficient Fault Tolerant NFS Server Designed for..](#) - Peyrouze, Muller (1996) (2 citations) (Correct)

....within the container. It is therefore impossible to use directly the Unix names of the files. **This led us to design a second file system at the user INRIA FT NFS: an Efficient Fault Tolerant NFS Server for Off the shelf Workstations 11 level and to separate file management from disk management [21], the latter being provided by Unix.** Instead of being internally denoted by an inode, a file in FT NFS is named with an ID. A first consequence of this choice is that file attributes are kept contiguously with the file data to allow the implementation of hardware links. **This has the important**

W. de Jonge, M.F. Kaashoek, and W.C. Hsieh. *The logical disk: A new approach to improving file systems*. In Proc. of 14th ACM Symposium on Operating Systems Principles, December 1993.

[Simplifying Distributed File Systems Using a Shared Logical Disk](#) - Shillner, Felten (1996) (6 citations) (Correct)

....view to users. **The result is a major simplification of the file system implementation. Our approach is based on picking some ideas from the work of others, and combining them with some new ideas, particularly in the area of fault tolerance. From de Jonge et al. we take the logical disk concept [8], which we extend to the distributed case.** From Attanasio et al. we take the idea of striping a logical device across distributed physical devices [2] which we extend by adding caching and cache coherence. **From Mann et al. we take the idea of expressing consistency requirements as ordering**

....extend by applying it to cache to cache transfers. **The result of combining these ideas, and adding some support for fault tolerance, is a base layer that makes it easy to write non centralized distributed file systems.** 1. 1 **Logical disk** **The shared logical disk is based on concept of logical disk [8].** A logical disk is an array of fixed size data blocks, which processes may use to stably store persistent data. **Processes** refer to the data blocks via logical block numbers; the logical disk maps these block numbers onto physical disk addresses. **The logical disk is free to store the blocks at any**

[Article contains additional citation context not shown here]

Wiedbren de Jonge, M. Frans Kaashoek, and Wilson C. Hsieh. *The logical disk: A new approach to improving file systems*. In Proceedings of the 14th Symposium on Operating System Principles, pages 15–28, 1993.

Symphony: An Integrated Multimedia File System - Shenoy, Goyal, Rao, Vin (1997) (27 citations) (Correct)

....the storage space is statically partitioned and the disk bandwidth is dynamically shared by logical volumes. In contrast, Symphony is a physically integrated file system, in which both storage space and disk bandwidth are dynamically shared among data types. Finally, the logical disk abstraction [15] provides an interface that allows multiple file systems to coexist on a single storage device. Logical disks provide functionalities similar to those provided by the data type independent layer of Symphony, such as multiple block sizes, location hints, etc. However, a key difference between

W. Jonge, M. F. Kaashoek, and W. C. Hsieh. *The Logical Disk: A New Approach to Improving File Systems*. In Proceedings of the Fourteenth Symposium on Operating Systems Principles, 1993.

A Comparison of Two Distributed Disk Systems - Lee, Thekkath, Whitaker, Hogg (1991) (1 citation) (Correct)

....sites because they only require access to a general purpose network and, by necessity, are designed to handle intermittent communication failures between nodes. 3. 2 Address Mapping Storage systems that manage multiple disks typically support virtual disk abstractions to simplify management [6, 19]. A virtual disk can span multiple physical disks for performance, mirror data for reliability, and can implement a variety of administrative policies. When clients access a virtual disk, the virtual disk addresses must be mapped to the corresponding physical disk addresses. Most storage systems

Wiebren de Jonge, M. Frans Kaashoek, and Wilson C. Hsieh. *The logical disk: A new approach to improving file systems*. In Proceedings of the ACM Symposium on Operating Systems Principles, pages 15–28, 1989.

A Comparison of Two Distributed Disk Systems - Lee, Thekkath, Whitaker, Hogg (1991) (1 citation) (Correct)

....and users become more sophisticated and increase beyond what can easily be satisfied by a few disk array controllers, managing the many disparate components of the storage system becomes a severe problem. There has been considerable research devoted to designing large scale storage systems [1 3, 5, 8, 10, 14, 18, 19]. A distributed disk system is a class of storage system that can reduce the complexity of building and managing large scale storage systems. Distributed disk systems manage collections of disks shared by, or partitioned across, multiple nodes as a single logical storage system that is highly

Wiebren de Jonge, M. Frans Kaashoek, and Wilson C. Hsieh. *The logical disk: A new approach to improving file systems*. In Proceedings of the 14th ACM Symposium on Operating Systems Principles, pages 15–28, December 1989.

A Case for Compositional File Systems - Bordawekar (Correct)

.... languages: Modula 3 or Java or C 6 Related Work The basic design approach uses the end to end argument [SRC84] The end to end argument has been used for designing operating and file systems [LS79, EKJ95] Other key projects that have influenced this design are: Exokernel Virtual disk [dJKH93] NASD file systems [GNA 97] xFS [ADN 96] Frangipani [TML97] HFS [Kri94] Forum [GSSW95] and Windows NTFS [Nag97, RvIG97] 7 Summary Traditional file systems suffer from the following disadvantages: application specific design, monolithic functional structure, and client server

W. de Jonge, M. Frans Kaashoek, and W. C. Hsieh. *The Logical Disk: A New Approach to Improving File Systems*. In Proceedings of the Fourteenth Symposium on Operating Systems Principles, December 1993.

MiSFIT: A Freely Available Tool for Building Safe Extensible Systems - Small (Correct)

....each test representing a class of possible OS extensions. We include a short description of each test; for more detail, the reader is directed to the earlier paper. hotlist: choose which page to evict from a linked list of page descriptors. lld: simulate the operation of a logical disk layer [DeJon93]. 6 . md5: compute the MD5 checksum [RFC1321] of 1MB of data. The tests were run on a 120MHz Pentium with 64MB of EDO memory, running BSD OS 2.1. Each test and its data easily fits into main memory. We report times relative to the unprotected version of the code. The results are found in Table

de Jonge, W., Kaashoek, M. F., Hsieh, W., "The Logical Disk: A New Approach to Improving File Systems,"

Proceedings of the 14th SOSP, 15–28, Asheville, NC (December 1993).

xFS: A Wide Area Mass Storage File System - Wang (1993) (16 citations) (Correct)

....table entry is of the following form: block ID # of blocks device address As data migrate among different storage levels, we simply change the corresponding translation tables, a simpler and cleaner approach than extending the existing UFS data structures. This is similar to the approach taken by [7] where logical disk addresses are mapped to physical ones to allow a clean separation between file and disk management without sacrificing performance. The memory management analogy, unfortunately, does not apply for data layout. Firstly, unlike processor caches which are usually direct mapped, we

W. de Jonge, M. F. Kaashoek, and W. Hsieh. *The Logical Disk: A New Approach to Improving File System Performance*. Proceedings of the 14th Symposium on Operating Systems Principles, December 1993. To appear.

The Swarm Scalable Storage System - John Hartman (1999) (8 citations) (Correct)

....as seen by a file system. data structures in the log. This is similar to running the LFS cleaner at user level, and enjoys the same advantages [15] Other examples of possible services are an atomic recovery unit (ARU) [6] service that provides atomicity across multiple log writes; a logical disk [4] service that provides a disk abstraction that hides the append only log, allowing higher level services and applications to overwrite the blocks they store; a caching service that caches log data in main memory; an encryption service; a compression service; etc. Distributed services, such as

W. De Jonge, M. F. Kaashoek, and W. C. Hsieh. *The Logical Disk: a new approach to improving file systems*. Operating Systems Review, 27(5):15–28, Dec. 1993.

Rodney Van Meter - Information Sciences (Correct)

....much like an extent based file system. Our work benefits from this work. Other possible file system structures, such as SGI's XFS [12] may depend on more dynamic, and hence complex, data structures, and may therefore not allocate blocks as predictably. A log based file system [9] or disk device [3] clearly will not, in their present forms, allocate blocks in a fashion amenable to improving throughput by careful choice of blocks. 6.2 Impact on File System Allocation Policies As proposed by Ghandeharizadeh [4] the idea of including a measure of ZCAV effects into a dynamic file relocater is

W. de Jonge, M. F. Kaashoek, and W. C. Hsieh. *The logical disk: A new approach to improving file systems*. In Proc. Fourteenth ACM Symposium on Operating Systems Principles, pages 15–28, Dec. 1993.

Elephant: The File System that Never Forgets - Douglas Santry (1999) (8 citations) (Correct)

.... a set of new utilities that exploit Elephant's novel functionality (we have already written a few utilities including tgrep , tls , and a history browser) Third, we are examining an alternate implementation that provides versioning at the level of blocks and abstracted by a logical disk [1]. Fourth, we are investigating how to backup an Elephant file system so that version histories can be recovered following a media failure. Finally, we are planning an extensive user study. To facilitate this study, we are modifying our prototype to allow it to shadow an NFS server. Users will thus

W. de Jonge, M. Kaashoek, and W. C. Hsieh. *The logical disk: A new approach to improving file systems*. In Proceedings of the 14th Symposium on Operating Systems Principles, pages 15–28, December 1993.

Opportunistic Log: Efficient Reads in a Reliable Object Server - O'Toole, Shrira (1994) (1 citation) (Correct)

....on the basis of disk head position, we can expect to dramatically reduce the expected cost of an installation read. There is substantial previous work on delayed processing of disk write operations [18, 10] Some methods applied to delayed disk writes involve writing pages at new locations [4, 7, 15] and would not work with disk reads. However, standard disk scheduling methods based on head position apply equally well to read operations. In particular, Seltzer et al. [18] have shown that when a pool of 1000 operations is available, greedy algorithms can reduce the cost of individual

W. de Jonge, F. Kaashoek, and W. Hsieh. *The logical disk: A new approach to improving file systems*. In Proc. of the 14th Symposium on Operating Systems Principles, Asheville, NC, December 1993. ACM.

Report of the Working Group on Storage I/O for Large-Scale .. - Gibson, Vitter, Wilkes (1996) (3 citations) (Correct)

.... onto spare disks, and future accesses remapped to the new devices [32] Storage devices are typically burdened by long positioning times, and a virtual device can be used to dynamically remap the physical location associated with a logical block, thus reducing the current access latency [19, 23, 57, 58, 73]. Additionally, most modern disk drives perform dynamic request reordering, in some cases taking advantage of low level information available only inside the storage device to optimize the request sequencing [35, 63] Since there is no single redundant disk array organization that is optimal for

de Jonge, W., Kaashoek, M. F., and Hsieh, W. C. *The logical disk: A new approach to improving file systems*. In Proc. of 14th ACM Symp. on Operating Systems Principles (December 1993).

The HP AutoRAID hierarchical storage system - Wilkes, al. (1996) (116 citations) (Correct)

....still further. Some of the schemes described in [Menon and Courtney 1993] are also used in the dual controller version of the HP AutoRAID array to handle controller failures. The Loge disk drive controller [English and Stepanov 1992] and its followons Mime [Chao et al. 1992] and Logical Disk [de Jonge et al. 1993], all used a scheme of keeping an indirection table to fixed sized blocks held on secondary storage. None of these supported multiple storage levels, and none was targeted at RAID arrays. Work on an Extended Function Controller at HP's disk divisions in the 1980s looked at several of these issues,

.... 1993; Seltzer et al. 1995] and cleaning (garbage collection) policies for them [Blackwell et al. 1995; McNutt 1994; Mogi and Kitsuregawa 1994] There is a large literature on hierarchical storage systems and the many commercial products in this domain (for example [Chen 1973; Cohen et al. 1989; DEC 1993; Deshpandee and Bunt 1988; Epoch Systems Inc. 1988; Gelb 1989; Henderson and Poston 1989; Katz et al. 1991; Miller 1991; Misra 1981; Sienknecht et al. 1994; Smith 1981] together with much of the proceedings of the IEEE Symposia on Mass Storage Systems) Most of this work has been concerned with

DE JONGE, W., KAASHOEK, M. F., AND HSIEH, W. C. 1993. *The Logical Disk: a new approach to improving file systems*. In Proceedings of 14th ACM Symposium on Operating Systems Principles. ACM, New York, 15--28.

A Comparison of OS Extension Technologies - Small, Seltzer (1996) (53 citations) (Correct)

....another region of the file. If the kernel uses heuristics (rather than application knowledge) to choose a read ahead policy, it can not cope with arbitrary application behavior. With the cooperation of the application, it can make more appropriate read ahead decisions. A Logical Disk facility (LD) [DEJON93] sits between the filesystem and the physical disk. The filesystem reads and writes logical blocks, and the LD maps the logical requests to locations on the physical disk. The LD can be used to transparently replicate data, by writing it in multiple places on the same disk or multiple disks, and

....to support a log structured layer between a filesystem and the physical disk. The simulation accepts write requests for logical blocks and maintains the mapping between these logical blocks and the physical blocks onto which they are stored. As with the system implemented by de Jonge et al. [DEJON93], our simulation maintains all data structures in main memory. We simulate a 1GB physical disk with 4KB blocks and 64KB (16 block) segments. Our simulation uses a stream of block write requests that are skewed so that 80 of the requests are for 20 of the blocks. Because our simulation does not

de Jonge, W., Kaashoek, M. F., Hsieh, W., "The Logical Disk: A New Approach to Improving File Systems," Proceedings of the 14th SOSP, pp. 15--28, Asheville, NC (December 1993).

DCD - Disk Caching Disk: A New Approach for Boosting I/O.. - Hu, Yang (1996) (6 citations) (Correct)

....approaches 80 . In addition, LFS needs to buffer a large amount of data for a relatively long period of time in order to write into disk later as a log, which may cause reliability problems. There are several other approaches such as log structured array [11] Loge [12] and Logical Disk approach [13]. The Logical Disk approach

improves the I/O performance by working at the interface between the file system and the disk subsystem. It separates file management from disk management by using logical block numbers and block lists. **Logical Disk** hides the details of disk block organization from the

W. de Jonge, M. F. Kaashoek, and W. C. Hsieh, "The logical disk: A new approach to improving file systems," in Proceedings of the 14th ACM Symposium on Operating Systems Principles, (Asheville, NC), pp. 15–28, Dec. 1993.

The HP AutoRAID hierarchical storage system - Wilkes, al. (1995) (116 citations) (Correct)

....mirrored form, thereby improving performance still further. Some of the schemes described in [Menon93] are also used in the dual controller version of the HP AutoRAID array to handle controller failures. **The Loge disk drive controller [English92] and its follow ons Mime [Chao92] and Logical Disk [deJonge93], all used a scheme of keeping an indirection table to fixed sized blocks held on secondary storage.** None of these supported multiple storage levels, and none were targeted at RAID arrays. Work on an Extended Function Controller at HP's disk divisions in the 1980s looked at several of these

Wiebren de Jonge, M. Frans Kaashoek, and Wilson C. Hsieh. *The Logical Disk: a new approach to improving file systems*. Proceedings of 14th ACM Symposium on Operating Systems Principles (Asheville, NC), pages 15–28, 5–8 December 1993.

Report of the Working Group on Storage I/O for Large-Scale .. - Gibson, Vitter, Wilkes (1996) (3 citations) (Correct)

.... onto spare disks, and future accesses remapped to the new devices [32] Storage devices are typically burdened by long positioning times, and a virtual device can be used to dynamically remap the physical location associated with a logical block, thus reducing the current access latency [19, 23, 57, 58, 73]. Additionally, most modern disk drives perform dynamic request reordering, in some cases taking advantage of low level information available only inside the storage device to optimize the request sequencing [35, 63] Since there is no single redundant disk array organization that is optimal for

de Jonge, W., Kaashoek, M. F., and Hsieh, W. C. *The logical disk: A new approach to improving file systems*. In Proc. of 14th ACM Symp. on Operating Systems Principles (December 1993).

Petal: Distributed Virtual Disks - Lee, Thekkath (1996) (133 citations) (Correct)

....work related to Petal in terms of four primary characteristics: type of abstraction (block level or file system level) degree of distribution, level of fault tolerance, and support for incremental expandability. **Related block level storage systems include RAID II [7] TickerTAIP [5] Logical Disk [8], Loge [10] Mime [6] AutoRAID [19] and Swift [4]** Some of these systems support only simple algorithmic mappings between the address space seen by a client and the underlying physical disks. This mapping is usually completely specified when the system is configured. In contrast, AutoRAID,

Wiebren de Jonge, M. Frans Kaashoek, and Wilson C. Hsieh. *The logical disk: A new approach to improving file systems*. In Proceedings of the 14th ACM Symposium on Operating Systems Principles, pages 15–28, December 1989.

Opportunistic Log: Efficient Installation Reads in a Reliable .. - O'Toole, Shrira (1994) (11 citations) (Correct)

....on the basis of disk head position, we can expect to dramatically reduce the expected cost of an installation read. **There is substantial previous work on delayed processing of disk write operations [13, 25] Some methods applied to delayed disk writes involve writing pages at new locations [5, 8, 22] and would not work with disk reads.** However, standard disk scheduling methods based on head position apply equally well to read operations. In particular, Seltzer et al. [25] have shown that when a pool of 1000 operations is available, greedy algorithms can reduce the cost of individual

W. de Jonge, F. Kaashoek, and W. Hsieh. *The Logical Disk: A New Approach to Improving File Systems*. In Proc. of the 14th Symposium on Operating Systems Principles, Asheville, NC, December 1993. ACM.

[Exodisk: Maximizing Application Control Over Storage Management - Grimm \(1996\)](#) (1 citation) **Self-citation (Kaashoek)** (Correct)

....to set their own storage allocation policies and optimizations. This approach prevents the exodisk system from transparently reorganizing data on disk or avoiding block overwrite semantics, as is done in some logical representations of disk storage (for example, in Mime [7] or the Logical Disk [8]) But it also does not suffer from the overheads associated with such logical to physical mappings, and it gives applications the flexibility to utilize the policies that are best suited for the given application. Furthermore, if desired, logical representations of disk storage can be implemented

....self protected extent, has its own protection information) Alternatively, it can use one exonode, storing the FAT in the exonode as application defined data and using fixed size extent entries for the individual blocks. **Logical representations of disk storage such as Mime [7] or the Logical Disk [8] map logical block numbers to physical disk locations.** Such a mapping can be efficiently implemented using an exonode: The logical block number can be an index into the array of exonode entries and the physical block information (again, a fixed size extent) is stored in that exonode entry.

[Article contains additional citation context not shown here]

Wiebren de Jonge, M. Frans Kaashoek, and Wilson C. Hsieh. *The logical disk: A new approach to improving file systems*. In Proceedings of the 14th Symposium on Operating Systems Principles, pages 15–28, Ashville, North Carolina, December 1993.

[Design Considerations for the Symphony Integrated.. - Shenoy, Goyal, Rao, Vin](#) (Correct)

No context found.

W. Jonge, M. F. Kaashoek, and W. C. Hsieh. *The Logical Disk: A New Approach to Improving File Systems*. In Proceedings of the Fourteenth Symposium on Operating Systems Principles, 1993.

[A Tool for Constructing Safe Extensible C++ Systems - Small \(1998\)](#) (7 citations) (Correct)

No context found.

de Jonge, W., Kaashoek, M. F., Hsieh, W., "The Logical Disk: A New Approach to Improving File Systems," Proc. 14th SOSP, Asheville, NC, 15– 28, December 1993.

[HFS: A flexible file system for shared-memory multiprocessors - Krieger \(1994\)](#) (17 citations) (Correct)

No context found.

Wiebren de Jonge, M. Frans Kaashoek, and Wilson C. Hsieh. *The logical disk: A new approach to improving file systems*. In Proceedings of the Fourteenth ACM Symposium on Operating Systems Principles, pages 15–28, 1993.

First 50 documents

[Online articles have much greater impact](#) [More about CiteSeer.IST](#) [Add search form to your site](#) [Submit documents](#) [Feedback](#)

CiteSeer.IST - Copyright [Penn State](#) and [NEC](#)